

BINARY DEPENDENCY DATABASE

Background of the Invention

5 Operating system configurations are composed of components that contain resources such as file and registry. Each component's footprint can be separated into two constituents: component resources and component dependencies. Footprint optimization faces many problems, including: packaging file/registry resources into components; satisfying required binary file dependencies; and modifying
10 existing large binaries into several smaller ones.

 Many operating systems targeted at embedded devices such as thin clients, retail point of sales, and set top boxes have limited space for storing the embedded OS, and the image size becomes crucial part of the final HW product cost. However, non-embedded devices, for which the image size is not a major issue, would
15 benefit from a smaller OS images as well. A smaller OS image uses less memory, boots faster, exposes a smaller hacker attack surface area, and reduces the likelihood of servicing the image.

 What is needed are tools and methods for analyzing binaries, components, configurations, and their footprints to help in component design.

20 **Summary of the Invention**

 The present invention is directed towards providing a foundation that facilitates the analysis of binaries, components, configurations, and their footprints to help in component design and optimization. Complete and meaningful binary, component, configuration, and footprint information helps to allow formal methods for
25 component analysis and configuration optimization.

 According to one aspect of the invention, a binary dependency database persists and stores binary dependency information. The binary dependency database

provides detailed dependency information among binaries. The binary dependency database allows dependencies to be linked that may span across binaries or functions.

According to another aspect of invention, a method and system are
5 directed to analyze binaries, components, and configurations and optimize their footprints meanwhile satisfying all their required dependencies.

According to yet another aspect of the invention, component verification and optimization tool are used by developers and testers for binary, component, and configuration footprint as well as dependency analysis. The optimization capability of
10 the tool allows users to analyze proposed changes to binaries, components, and configurations in order to improve footprints against given constraints.

Brief Description of the Drawings

FIGURE 1 illustrates an exemplary computing device that may be used in one exemplary embodiment of the present invention;
15 FIGURE 2 illustrates a first order dependency vector;
FIGURE 3 illustrates a first order dependency matrix;
FIGURE 4 shows exemplary 2nd, 3rd and 4th order dependency matrices;
FIGURE 5 illustrates a flow for creating the binary dependency database;
FIGURE 6 shows exemplary types of dependencies that a binary may be
20 dependent on; and
FIGURE 7 illustrates exemplary information that is maintained by the binary dependency database, in accordance with aspects of the invention

Detailed Description of the Preferred Embodiment

Illustrative Operating Environment

25 With reference to FIGURE 1, one exemplary system for implementing the invention includes a computing device, such as computing device 100. In a very basic configuration, computing device 100 typically includes at least one processing unit 102 and system memory 104. Depending on the exact configuration and type of

computing device, system memory 104 may be volatile (such as RAM), non-volatile (such as ROM, flash memory, etc.) or some combination of the two. System memory 104 typically includes an operating system 105, one or more applications 106, and may include program data 107. In one embodiment, application 106 may include application
5 120 relating to a binary dependency database. This basic configuration is illustrated in FIGURE 1 by those components within dashed line 108.

Computing device 100 may have additional features or functionality. For example, computing device 100 may also include additional data storage devices (removable and/or non-removable) such as, for example, magnetic disks, optical disks,
10 or tape. Such additional storage is illustrated in FIGURE 1 by removable storage 109 and non-removable storage 110. Computer storage media may include volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information, such as computer readable instructions, data structures, program modules, or other data. System memory 104, removable
15 storage 109 and non-removable storage 110 are all examples of computer storage media. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired
20 information and which can be accessed by computing device 100. Any such computer storage media may be part of device 100. Computing device 100 may also have input device(s) 112 such as keyboard, mouse, pen, voice input device, touch input device, etc. Output device(s) 114 such as a display, speakers, printer, etc. may also be included. These devices are well known in the art and need not be discussed at length here.

25 Computing device 100 may also contain communication connections 116 that allow the device to communicate with other computing devices 118, such as over a network. Communication connection 116 is one example of communication media. Communication media may typically be embodied by computer readable instructions, data structures, program modules, or other data in a modulated data signal, such as a
30 carrier wave or other transport mechanism, and includes any information delivery

media. The term “modulated data signal” means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as
5 acoustic, RF, infrared and other wireless media. The term computer readable media as used herein includes both storage media and communication media.

Binary Dependency Database

The present invention is directed to facilitate the analysis of binaries, components, and configurations by providing complete and meaningful information that
10 will aid in component analysis and configuration optimization.

Table 1 illustrates an exemplary operating system repository that contains 11,069 files with total size of approximately 0.99 GB. As can be seen from Table 1, three specific file types, DLL, EXE and SYS, are responsible for the half of the repository size. There are 2943 (27%) repository files that have 17,849 dependencies of
15 the following types: static; forward reference image; bound; delay load version 1; and delay load version 2.

In one exemplary repository it was found that approximately 95 of these dependencies are due to the three file types: DLL, EXE and SYS. There may also be dynamic dependencies among repository files that are created by using an API. For
20 example, the following Win32 APIs may be used that create dynamic dependencies, including: LoadLibrary; CreateProcess; and CoCreateInstance. Other operating systems include other calls that create dynamic dependencies.

Table 1: Repository dissection by different file types

File type	Total number of files	Total repository file size [MB]	Repository size ratio [%]	Average file size [kB]
DLL	2176	386	39.0%	182
TTC	6	63.4	6.4%	10,820
SYS	663	61.7	6.2%	95
EXE	452	47.4	4.8%	107
ICM	243	40.2	4.1%	169
DIC	11	36.4	3.7%	3,389
GPD	1826	35.1	3.5%	20
PPD	939	29.6	3.0%	32
CHM	267	26.5	2.7%	102
TTF	127	23.8	2.4%	192
INF	691	21.3	2.2%	32

Table 2 summarizes the dependency type breakdown of the dependencies listed in Table 1. Row#1 can be read as follows – there are 936 files in the example repository that have total of 2895 static dependencies on 135 distinct files; they account for 16.2% of all dependencies; all these dependencies are required.

Table 2: Binary file dependency types

# of files	Dependency type	# of deps	Dependencies on # of distinct files	Fraction of all deps [%]	Required dependency
936	Static	2895	135	16.2	YES
7	Bound	7	1	0.04	YES
2009	Static + Bound	12425	368	69.6	YES
1277	Bound + Forward reference image	1350	11	7.6	YES
462	Static + Bound + Forward reference image	479	8	2.7	YES
139	Delay load version 2	504	119	2.8	MAYBE
26	Delay load version 1	189	40	1.1	MAYBE

Note: One file can have more than one dependency type.

- All dependency types containing “static” or “bound” dependencies are required dependencies with respect to Windows OS, and they are responsible for 96.1% of all dependencies. Delay load dependencies can be either optional or required and are responsible for only 3.9% of all dependencies.

Binary Dependency Type, Order and Strength

- In some cases one binary will not relate to another just through single dependency type, order, or strength, and the relation will be more complex between two files in the system.

- Let’s consider first order dependencies. A first order dependency is a direct dependency relationship. For example, consider the following. Certmgr.dll directly depends on netapi32.dll through a static dependency, and it is importing only 4 functions out of a possible 331 functions exported by netapi32.dll. The strength of the

bond between these two binaries is assessed to be 4/331, which is ~1.2%. However, certmgr.dll has also static dependencies on

certcli.dll (strength = 16/141 ~ 11.3%)

shell32.dll (strength = 4/754 ~ 0.5%)

5 cryptui.dll (strength = 9/48 ~ 18.8%)

ntdsapi.dll (strength = 4/96 ~ 4.2%)

All four of these files again directly depend on netapi32.dll. Therefore, certmgr.dll has also four second order dependencies on netapi32.dll. Moreover, certmgr.dll has also third and fourth order dependencies on netapi32.dll such as

10 certmgr.dll → mfc42u.dll → winspool.drv → netapi32.dll

certmgr.dll → advapi32.dll → secur32.dll → netapi32.dll

certmgr.dll → oleaut32.dll → advapi32.dll → secur32.dll → netapi32.dll

certmgr.dll → wintrust.dll → advapi32.dll → secur32.dll → netapi32.dll,

wherein “→” means depends on.

15 Since there are 176 repository files directly depending on netapi32.dll (150 through static and 26 through delay-load dependencies) the dependency of certmgr.dll on netapi32.dll will most probably grow through several more orders.

This example illustrates that dependency relations between two files can be more complicated than just a direct dependency. The dependency requirement, such as required vs. optional, will also be different along different dependency paths since
20 some of the dependencies along the path are required (such as static) and some of them might be optional (such as delay loads). Given this complexity, it is useful to provide to developers the possible dependency paths between two specified binaries.

More dependency data regarding order, strength, and common functions
25 called with a common dependency may also be generated to identify dependency orders, types, strengths, and any common functions between two binaries among all possible paths. Information like this will be used to determine the strength of the coupling between two binaries.

30

Binary, Component, and Configuration Analysis

- When analyzing existing binaries, components, and configurations or when making a new binary or component dependency it is useful to know what the dependency and footprint ramifications will be and what specific binaries, components and configurations will be affected throughout the dependency chain. Such information is not only useful with regards to the footprint but is also directly related to the functional testing of binaries, components and configurations.

Binary Dependency Chain Analysis

- Table 3 illustrates exemplary binary dependency tree footprints as well as exemplary dependency. While only a number of dependencies and their total size is specified in this table, all dependencies, their sizes and footprint ratios are calculated to help identify the largest footprint files within the dependency chain.

Table 3: Binary dependency tree analysis

Binary	Binary size [kB]	Binary component owner	# of binary deps.	Total binary dep footprint [kB]	Binary/Total footprint [%]
Moviemk.exe	789	Windows Movie Maker	121	48,139	1.6
Wmpcore.dll	1,272	Windows Media Player 8.0	119	47,084	2.7
Wmplayer.exe	508	Windows Media Player 8.0	118	47,599	1.1
Msrating.dll	129	Internet Explorer	109	43,588	0.3
Srv.sys	323	File Sharing	9	2,633	12.3
Csrsrv.dll	29	Client/Server Runtime (Console)	2	687	4.2

- Therefore, when making a new dependency, e.g. to “msrating.dll”, then there is a dependency on 109 binaries with total size of ~43 MB. However,

“msrating.dll” belongs to the “Internet Explorer” component that is a fairly large component, 4.2 MB component itself and ~63 MB with its dependencies.

It is also useful to know what other binaries depend directly on a particular binary and throughout the whole dependency chain, and what components these binaries belong to. Such information is useful in determining how the binaries, components and their footprints will be affected by a code or dependency change in the binary in question.

In Table 4 below, three binaries on which most other binaries depend throughout the whole dependency chain are listed, as well as three randomly chosen binaries. According to one embodiment, binary sizes, components which own the depending binaries, and the component sizes are also displayed.

Table 4: Reverse binary dependency tree analysis

Binary	Component Owner	# of binaries directly depending on the binary	Total # of binaries throughout the whole dependency chain depending on the binary
Ntdll.dll	Primitive: Ntdll	1755	2219
Kernel32.dll	Win32 API – Kernel	2168	2203
Msvcrt.dll	Microsoft Visual C++ Run Time	1605	2121
Clbcatq.dll	Primitive: Clbcatq	6	15
Comres.dll	Primitive: Comres	20	36
Csrsrv.dll	Client / Server Runtime (Console)	4	4

It might be obvious to the owner of kernel32.dll that any change in footprint or dependency in this file will affect the whole operating system, but it might not be obvious to the owner of clbcatq.dll that a change in this binary will affect six (6) other binaries directly and total of fifteen (15) binaries both directly and indirectly.

Analysis raises awareness among binary and component owners of their binary effect on other binaries and components throughout the system, as well as helping testers to identify areas of testing once there is a change in a particular binary.

- Orphan binaries: These are binaries that do not depend on any other
 5 binaries and no other binaries depend on them.

Binary dependency clusters: A cluster is a set of binary files that depend on each other in a circular way virtually creating one large binary.

- File resource packaging in components: Most components contain more than one file resource and it is useful to look at how to choose these resources for given
 10 components knowing the binary dependencies. Table 7 illustrates exemplary file resources of three components and their dependency tree sizes.

Table 7: Component file resources and their dependency tree footprints

Component Name	File Resources	File sizes [kb]	# of deps	Size of dep chain [kb]	# of common deps	Binaries depending on file resources
Net.exe Utility	Net.exe	39	110	43,626	None	None
	Net1.exe	113	110	43,700		
	Net.hlp	101	1	0		
	Neth.dll	248	1	0		
TCP/IP Networking	Wshtcpip.dll	17	110	43,605	None	None
	Tcpip.chm	50	1	0		
	Tcpip.sys	320	7	2,547		
Client / Server Runtime (Console)	Winsrv.dll	270	112	43,930	Csrsvr.dll Ntdll.dll	Basesrv.dll, csrss.exe, winsrv.dll depend on csrssrv.dll
	Csrss.exe	4	3	691		
	Csrsvr.dll	29	2	687		

Generally, help files in these components, such as *.hlp, *.chm, and *.cnt, are not much of interest since these files can be suppressed from being copied during building an image so they can be considered as optional. However, there are binaries that will not function properly when their help files are not present.

- 5 Configuration analysis: Many configurations will share many components due to the common dependencies; therefore, footprint of two configurations is not sum of their footprints, as it is in components only footprints, but is a union of the two configurations.

- 10 It is valuable to know what components and files compose specific configurations and what the footprint, component, and binary ramifications will be once two or more configurations are combined. There are two different scenarios shown in Table 8.

Table 8: Configuration-Component comparison

Configuration	Size [MB]	# of components	# of common components	# of exclusive components
Minlogon	10.5	38	36	Primitive: Userenv Minlogon
Command Shell	9.9	37		Command Shell
Command Shell	9.9	37	36	Command Shell
Explorer Shell	68.4	186		150 components
Config containing Notepad.exe	72.6	186	146	40 components
Explorer Shell	68.4	186		

For example, a user wants to know what the Minlogon configuration footprint is, what components and files it contains, and what the ramifications will be once “Command Shell” configuration will be chosen for Minlogon. As Table 8 shows, these two configurations have 36 out of possible 38

5 components in common, and the footprint difference between Minlogon and Command Shell would be 0.6 MB. Minlogon configuration in fact has all the components required by Command Shell, except the Command Shell component itself. Moreover, if there is a need to add Notepad.exe to Minlogon configuration, the footprint will increase from ~10 MB to ~73 MB.

10 As a second example, one may be interested in choosing between Command Shell and Explorer Shell. From Table 8, it can be seen that the footprint difference, 55.5 MB, and the difference in number of components used by Command and Explorer shell configurations is significant.

15 It is interesting to note that five largest components in Minlogon and Command Shell configuration, out of all 38 and 37, respectively, are responsible for over 50% of total configuration footprints. In the “Explorer Shell” configuration, it is the largest 9 components, out of total 186, that are responsible for ~50% of total footprint, as can be seen in Table 9.

Table 9: Details of configuration footprints

Configuration	Largest components within configuration	Size [bytes]	Ratio [%]
Minlogon	---	11,048,945	---
	Win32 API - User	2,360,704	21.4%
	Win32 API - GDI	969,088	8.8%
	Win32 API - Kernel	926,720	8.4%
	NLS: Core Files	728,733	6.6%
	Primitive: Ntdll	674,304	6.1%
Command Shell	---	10,391,163	---
	Win32 API - User	2,360,704	22.7%
	Win32 API - GDI	969,088	9.3%
	Win32 API - Kernel	926,720	8.9%
	NLS: Core Files	728,733	7.0%
	Primitive: Ntdll	674,304	6.5%
Explorer Shell	---	68,528,592	---
	Winlogon Support	13,127,585	19.2%
	Primitive: Shell32	8,322,560	12.1%
	User Interface Core	2,849,069	4.2%
	Win32 API - User	2,360,704	3.4%
	Microsoft Foundation Class Library (MFC)	1,990,767	2.9%
	Windows Logon (Standard)	1,575,936	2.3%
	Local Security Authority Subsystem (LSASS)	1,414,192	2.1%
	Primitive: Shdocvw	1,338,880	2.0%
	Primitive: Ole32	1,141,248	1.7%

The basis for formal methods is complete and accurate information. This means thorough knowledge of the relationships between all binaries that comprise

software. According to one embodiment, this information is captured, persisted, updated, refined, and shared via a binary dependency database.

Binary, Component and Configuration Analysis Tool

FIGURES 2-4 illustrate identifying binary clusters and binaries related to clusters, by identifying common and exclusive binaries in multiple different dependency trees and identifying binaries depending on a specific binary through second (third, fourth, and so forth) order dependency.

FIGURE 2 illustrates a first order dependency vector, in accordance with aspects of the invention. According to one embodiment of the invention, matrices for binary, component, and configuration analysis and their dependencies are used. Matrices have been used in analyzing and optimizing various types of problems such as marriage problem, transportation problems, spanning trees, and various network problems.

Consider the following vector (1,0,0,0,0,1,0) representing specific binary (i.e. kernel32.dll) direct dependencies, 1s representing dependencies and 0s representing non-dependencies.

Looking at the 0s and 1s in the vector, it can be identified that kernel32.dll depends on kernel32.dll and ntdll.dll and no other binaries. Once the same vector representation for the remaining seven binaries is performed, a square matrix can be generated as shown in FIGURE 3.

FIGURE 3 illustrates a first order dependency matrix, in accordance with aspects of the invention. The matrix provides additional information when its column vectors (binaries that depend on a specific binary) are examined.

Turning to FIGURE 3 one can see that every single binary represented in this matrix depends on ntdll.dll since the column vector for ntdll.dll contains all 1s. The same information can be obtained by looking at the column vector of the specific binaries.

The elements of this dependency matrix are describing direct dependencies (i.e. first order dependencies) only. There are different ways to get n^{th}

order dependency matrix. Determine each binary's n^{th} order dependencies and populate the matrix rows appropriately or multiply the matrix $(n-1)$ times.

Since 0s and non-zeros are looked at according to one embodiment, it is not necessary to multiply the matrix in the mathematical sense but abort computation
5 once the entry for a matrix element is or has reached a non-zero value.

FIGURE 4 shows exemplary 2^{nd} , 3^{rd} and 4^{th} dependency order matrices, in accordance with aspects of the invention.

Note the matrix elements in bold font; these are the elements that differ from the previous dependency order matrix elements. After the 5^{th} order dependency
10 matrix was generated (not shown) no zero elements changed from the previous 4^{th} order dependency matrix. No changes indicate that the 4^{th} order dependency matrix is the final full dependency matrix. The rows of this matrix represent the full dependency lists of tree elements of the specific binaries representing the specific rows.

Further introducing a file size column vector $S = (905, 1115, 633, 793,$
15 $990, 659)$ that represents sizes in kB of the binaries in the above matrix, one can determine the size of the dependency trees for all binaries in the above matrix by simply multiplying the above dependency matrix by the vector S ; therefore $D' \cdot S$ (where D' is the reduced dependency matrix) will give us the sizes of all binary dependency trees for the binaries specified by rows of matrix D' . Note that the vector S size matches the size
20 of D' in order for the inner product to make sense. D' can be any order dependency matrix and it can be used to compare how binary dependency tree footprints degrade once we go deeper into dependencies order (i.e. one can compare how footprint of the 1^{st} order binary dependency trees compare to the full order dependency trees).

The items discussed so far are just simple analysis of given designs such
25 as given binaries, components and configurations. Optimization of the design discussed will now be discussed.

Binary, Component and Configuration Optimization Tool

Optimization results in the best design given certain constraints. Since there are typically thousands of binaries in the repository and about the same amount of
30 components in the component database, it is not trivial to produce an optimal design. In

order to proceed with the optimization work several types of constraints are identified. According to one embodiment, these include: number of binaries in the system; binary sizes; number of components in the database; component sizes; and configuration footprints.

5 Binary Dependency Database (BDD)

FIGURE 5 illustrates a flow for creating the binary dependency database, in accordance with aspects of the invention. As illustrated, FIGURE 5 includes binary dependency database (BDD) 505, binaries 510, dependency tool(s) 515, registry dependencies 516, file, font, icon dependencies 517, binary dependencies 518, functional dependencies 519, other dependencies 520, database importing and
10 dependency resolution tool 525, web interface 545, component analyzer 550, analyzer 535, and user input 555.

BDD 505 receives information relating to dependencies 516-520 through tool 525. Dependency tool(s) are any tools that capture dependencies relating to
15 binaries 510 and/or source code 530. Users may also input dependency information through user input 555 accessible using web interface 545. Dependency information may also be obtained during runtime of the binaries. Component analyzer 550 may be used by developers to analyze configuration designs and footprint analysis.

BDD 505 captures and persists the relationships between binaries 510
20 and provides a readily accessible dependency model of a software system.

According to one embodiment, binary dependency database 505 is separate and distinct from a component database (not shown) that ships with an operating system, such as Windows XPE. The component database represents an abstract, macro view of the OS and is defined in terms of components. Binary
25 dependency database 505, in contrast, is a fundamental, micro view of the OS and is defined in terms of binary files and functions.

BDD 505 captures the relationships between the executable files (binaries 510) that comprise a software system. According to one embodiment, for every executable, the BDD contains the names of the files on which the executable
30 depends but also detailed attributes of each dependency. The BDD also includes

functions, as well as other files, upon which the executable depends. According to one embodiment of the invention, the BDD is stored in a SQL database. The BDD links dependency information between binaries that may span across binaries and functions. According to one embodiment, these dependency attributes include:

- 5 Dependency Category - Identifies the dependency relationships such as static, delay-load, dynamic (LoadLibrary, CoCreateInstance, CreateProcess), or registry.

Static dependencies are established when the executable is built and can be easily identified by inspecting the binary file. For example, if A.EXE statically links
10 to B.DLL, then A.EXE is statically dependant on B.DLL. The static dependency is readily identified by inspecting the contents of A.EXE.

Delay-load relationships are a special kind of static dependency. They too are easily identified by inspecting the contents of A.EXE.

Dynamic dependencies are established at runtime and are more difficult
15 to identify. For example, if A.EXE calls LoadLibrary("B.DLL"), then A.EXE has a dynamic dependency on B.DLL.

A binary is considered dependent on a registry entry if it opens or appends a registry key, or queries or appends a registry value.

Dependencies 516-520 include many different dependency types upon
20 which the binaries and functions within which the binaries may depend. According to one embodiment the dependencies are identified as optional or required. It is difficult to distinguish between optional and required dependencies. Whether the dependant is optional or required is determined by design rather than implementation. Therefore, the dependency type typically requires input from developers and program managers who
25 can consult the appropriate design specification.

Dependency Strength is used to quantify the strength of the bond between two binaries. For example, if A.EXE imports two of the ten functions exported by B.DLL, then the dependency strength between A.EXE and B.DLL is 2/10, or 0.2.

BDD 505 is populated with data generated by the database importing and dependency resolution tool 525. The BDD is populated using static and dynamic dependency information and may be viewed using a web-based interface (545).

Binary Dependency Database Supporting Tools

5 Dependency tool(s) 515 output the static and dynamic, dependency information. According to one embodiment, tool 515 outputs not only the files that the binary depends on, but also the functions and registry data on which it depends.

 Web-based database interface 545 has read/write access to binary dependency database 505. Analyzer 535 and component analyzer 550 are used to
10 analyze data from BDD 505. According to one embodiment, the data output includes: Binaries dependent on the specified binary(s); Binaries the specific binary(s) depend on; Binary file and component 1st, 2nd, 3rd... nth order dependencies; the functions, files, registry data, and other dependencies within the system that the binary is dependent on; and the footprint growth comparison between different dependency orders taken into
15 the account; Footprint comparisons and enumeration of optional vs. required binary dependencies; Orphan binaries; Unused DLL export functions; Component(s) and configuration(s) specified binary(s) belong to; Common and exclusive dependencies among specified binaries; Common and exclusive components across multiple configurations; Dependency strength between specified binaries; and Relative footprint
20 ratio for binaries, components, and configurations.

 FIGURE 6 shows exemplary types of dependencies that a binary may be dependent on, in accordance with aspects of the invention. As illustrated in FIGURE 6, dependency information for binary 605 may be registry dependency information 610,
file, font, icon information 615, binary information 620, function information 625, and
25 other information 630.

 The binary dependency database may store dependency information related to many different types of files. Registries can include dependencies relating to binaries. For example, a binary or binaries may depend on a key being contained within the registry. This dependency may be stored within the BDD. A binary may also
30 depend on a specific file, icon, font or some other material being present. For example,

a binary may have a dependency on a non-standard font. Binary 605 may also be dependent on binaries (620) and functions (625). Binary 605 may also be dependent on other information.

FIGURE 7 illustrates exemplary information that is maintained by the binary dependency database, in accordance with aspects of the invention. As illustrated, BDD 705 includes source and destination names, source and destination paths, dependency information, function information, size information, and type of file information.

A file may be known by many different names. For example, a file may have a source name on a distribution that is different than its name after it is installed (the destination name). Similarly, files may have the same name but be located under different paths on an installed image. For example, different DLLs may have the same name but be located at different paths. Without keeping track of the path information it would be difficult to know which file the binary is actually dependent on. Function information relating to the binaries is also maintained. For example, binary A may depend on functions F1, F2, F3, and F4 wherein F4 depends upon F2 located within Binary B. Maintaining the function information in addition to the binary dependencies further refines the dependency information. As a result, a developer could possibly include only a function from a binary instead of including the entire binary within a build.

The above specification, examples and data provide a complete description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made without departing from the spirit and scope of the invention, the invention resides in the claims hereinafter appended.